



From low-rank retractions to dynamical low-rank approximation and back

Axel Séguin¹ · Gianluca Ceruti² · Daniel Kressner¹

Received: 12 September 2023 / Accepted: 4 June 2024

© The Author(s) 2024

Abstract

In algorithms for solving optimization problems constrained to a smooth manifold, retractions are a well-established tool to ensure that the iterates stay on the manifold. More recently, it has been demonstrated that retractions are a useful concept for other computational tasks on manifold as well, including interpolation tasks. In this work, we consider the application of retractions to the numerical integration of differential equations on fixed-rank matrix manifolds. This is closely related to dynamical low-rank approximation (DLRA) techniques. In fact, any retraction leads to a numerical integrator and, vice versa, certain DLRA techniques bear a direct relation with retractions. As an example for the latter, we introduce a new retraction, called KLS retraction, that is derived from the so-called unconventional integrator for DLRA. We also illustrate how retractions can be used to recover known DLRA techniques and to design new ones. In particular, this work introduces two novel numerical integration schemes that apply to differential equations on general manifolds: the accelerated forward Euler (AFE) method and the Projected Ralston–Hermite (PRH) method. Both methods build on retractions by using them as a tool for approximating curves on manifolds. The two methods are proven to have local truncation error of order three. Numerical experiments on classical DLRA examples highlight the advantages and shortcomings of these new methods.

Communicated by n/a.

Parts of this work were performed while Gianluca Ceruti was part of the EPFL Institute of Mathematics, and it was supported by the SNSF research project *Fast algorithms from low-rank updates*, grant number: 200020_178806.

✉ Daniel Kressner
daniel.kressner@epfl.ch

Axel Séguin
axel.seguin@alumni.epfl.ch

Gianluca Ceruti
gianluca.ceruti@uibk.ac.at

¹ École Polytechnique Fédérale de Lausanne (EPFL) Institute of Mathematics, 1015 Lausanne, Switzerland

² University of Innsbruck, Numerical Analysis and Scientific Computing, Innsbruck, Austria

Keywords Dynamical low-rank approximation · Retraction · Time integration

Mathematics Subject Classification 65L05 · 65L20

1 Introduction

This work is concerned with the numerical integration of manifold-constrained ordinary differential equations (ODEs) and, in particular, the class of ODEs evolving on fixed-rank matrix manifolds encountered in dynamical low-rank approximation (DLRA) [20]. This model-order reduction technique has attracted significant attention in recent years, finding numerous applications in fundamental sciences and engineering [10, 16, 22, 25]. A key feature of DLRA is that it can bring down the computational cost, in terms of storage and computational time, without compromising accuracy when solving certain large-scale ODEs as they arise, for instance, from the discretization of partial differential equations.

Several works have already pointed out the close connection of specific DLRA techniques to retractions on fixed-rank manifolds, the latter featuring prominently in Riemannian optimization algorithms [1, 5]. For instance, the projector-splitting integrator for DLRA proposed in [24] is shown in [4] to define a second-order retraction. The other way around, the more frequently encountered fixed-rank retraction defined with the metric projection on the manifold is centrally used in the projected Runge–Kutta methods for DLRA developed in [19]. Furthermore, the recently proposed dynamically orthogonal Runge–Kutta [9] schemes extend Runge–Kutta schemes to the DLRA setting by making use of high-order approximations to the metric projection retraction. In this work, we consider a unified framework for relating popular DLRA time integration techniques to particular retractions on the fixed-rank manifold. This not only covers the existing relations mentioned above but it also allows us to provide a novel geometric interpretation of the so-called unconventional integrator [8] by showing that it defines a second-order retraction that coincides with the orthographic retraction modulo a third-order correction term.

Low-rank retractions can also be used to design new time integration algorithms on fixed-rank matrix manifolds, based on the insight [29] that retractions can be used to conveniently generate manifold curves. For instance, any retraction allows one to define a manifold curve passing through a prescribed point with prescribed velocity. If the retraction additionally possesses the so-called second-order property, the acceleration can also be prescribed. In addition, as put forth in [26], a retraction can be used to define a smooth manifold curve with prescribed endpoints and prescribed endpoint velocities. We use these two retraction-based curves as the building block for two novel numerical integration schemes for manifold-constrained ODEs and in particular for DLRA, namely the accelerated forward Euler method and the Projected Ralston–Hermite (PRH) method.

The proposed novel numerical integrators are then tested in two classical scenarios: the Lyapunov differential equation and a synthetic example [18, §2.1]. The first scenario serves as an effective numerical test to examine the impact of model errors, i.e., the magnitude of the normal component of the projected vector field, while the second

analyzes the proposed method's robustness in dealing with small singular values, a notable concern in the context of numerical integrators for DLRA. We showcase that the AFE numerical method maintains its robustness when dealing with small singular values but is sensitive to the presence of significant model errors. In contrast, the PRH method is more susceptible to the influence of small singular values but maintains accuracy levels comparable to the second-order Projected Runge–Kutta (PRK) scheme [19]. Furthermore, for large model errors, it is illustrated that the computational time of the proposed PRH method is comparable to that of the PRK counterpart. The computational time of the accelerated methods is primarily dominated by the computation of the acceleration, namely, the Weingarten map.

Outline Sect. 2 collects preliminary material on differential geometry and DLRA. The link between retractions and existing DLRA techniques is discussed in Sect. 3. Then, Sect. 4 presents two new retraction-based integration schemes: the Accelerated Forward Euler method in Sect. 4.1 and the Projected Ralston–Hermite integration scheme in Sect. 4.2. Numerical experiments reported in Sect. 5 assess the accuracy and stability to small singular values of both methods and provide performance comparison with state-of-the-art techniques. In Sect. 6, a discussion on future work directions concludes the manuscript.

2 Preliminaries

2.1 Differential geometry

In this section, we briefly introduce concepts from differential geometry needed in this work. The reader is referred to [5] for details.

In the following, \mathcal{M} denotes a manifold embedded into a finite-dimensional Euclidean space \mathcal{E} . When necessary, \mathcal{M} is endowed with the Riemannian submanifold geometry: for every $x \in \mathcal{M}$, the tangent space $T_x\mathcal{M}$ is equipped with the induced metric inherited from the inner product of \mathcal{E} . The orthogonal complement of $T_x\mathcal{M}$ in \mathcal{E} is the normal space denoted by $N_x\mathcal{M}$. The orthogonal projections from \mathcal{E} to $T_x\mathcal{M}$ and $N_x\mathcal{M}$ are respectively denoted by $\Pi(x)$ and $\Pi^\perp(x) := (I - \Pi)(x)$. The disjoint union of all tangent spaces is called the tangent bundle and denoted by $T\mathcal{M}$. Every ambient space point $p \in \mathcal{E}$ that is sufficiently close to the manifold admits a unique metric projection onto the manifold [3, Lemma 3.1] and we denote it by $\Pi_{\mathcal{M}}(p) := \arg \min_{x \in \mathcal{M}} \|p - x\|$.

For a smooth manifold curve $\gamma : \mathbb{R} \rightarrow \mathcal{M}$, its tangent vector or velocity vector at $t \in \mathbb{R}$ is denoted by $\dot{\gamma}(t) \in T_{\gamma(t)}\mathcal{M}$. If the curve γ is interpreted as a curve in the ambient space \mathcal{E} , or if γ maps to a vector space, we use the symbol $\gamma'(t)$ to indicate its tangent vector at $t \in \mathbb{R}$.

The Riemannian submanifold geometry of \mathcal{M} is complemented with the Levi-Civita connection ∇ which determines the intrinsic acceleration of a smooth manifold curve γ as $\ddot{\gamma}(t) := \nabla_{\dot{\gamma}(t)}\dot{\gamma}(t)$, $\forall t \in \mathbb{R}$. We reserve the symbol $\gamma''(t)$ to indicate the extrinsic acceleration of the curve intended as a curve in the ambient space \mathcal{E} .

Manifold of rank- r matrices. The leading example in this work is the case where $\mathcal{E} = \mathbb{R}^{m \times n}$ and $\mathcal{M} = \mathcal{M}_r$, the manifold of rank- r matrices. We represent rank- r

matrices and tangent vectors to \mathcal{M}_r with the conventions used in [31] as well as in the MATLAB library Manopt [6] that was used to perform the numerical experiments. Any $Y \in \mathcal{M}_r$ is represented in the compact factored form $Y = U \Sigma V^\top \in \mathcal{M}_r$ with $U \in \mathbb{R}^{m \times r}$, $V \in \mathbb{R}^{n \times r}$ both column orthonormal and $\Sigma = \text{diag}(\sigma_1, \dots, \sigma_r)$ containing the singular values ordered as $\sigma_1 \geq \dots \geq \sigma_r$. Any tangent vector $T \in T_Y \mathcal{M}_r$ is represented as $T = U M V^\top + U_p V^\top + U V_p^\top$, where $M \in \mathbb{R}^{r \times r}$, $U_p \in \mathbb{R}^{m \times r}$ and $V_p \in \mathbb{R}^{n \times r}$ such that $U^\top U_p = V^\top V_p = 0$. The manifold \mathcal{M}_r inherits the Euclidean metric from $\mathbb{R}^{m \times n}$ so each tangent space is endowed with the inner product $\langle W, T \rangle = \text{Tr}(W^\top T)$ and the Frobenius norm $\|W\| = \|W\|_F$, for any $W, T \in T_Y \mathcal{M}_r$. Then, the orthogonal projection onto the tangent space at Y of any $Z \in \mathbb{R}^{m \times n}$ is given by

$$\Pi(Y)Z = UU^\top ZVV^\top + (I - UU^\top)ZVV^\top + UU^\top Z(I - VV^\top). \tag{2.1}$$

The metric projection onto the manifold \mathcal{M}_r of a given $A \in \mathbb{R}^{m \times n}$ is uniquely defined when $\sigma_r(A) > \sigma_{r+1}(A)$. It is computed as the rank- r truncated SVD of A and we denote it by $\Pi_{\mathcal{M}_r}(A)$.

Retractions. A retraction is a smooth map $R : T\mathcal{M} \rightarrow \mathcal{M} : (x, v) \mapsto R_x(v)$ defined in the neighborhood of the origin x of each tangent space $T_x \mathcal{M}$ such that it holds that $R_x(0) = x$ and $DR_x(0)[v] = v$, for all $v \in T_x \mathcal{M}$. The defining properties of a retraction can be condensed into the fact that for any $x \in \mathcal{M}$ and $v \in T_x \mathcal{M}$ the map $\tau \in \mathbb{R} \mapsto \sigma_{x,v}(\tau) = R_x(\tau v)$ is well-defined for any sufficiently small τ and parametrizes a smooth manifold curve that satisfies $\sigma_{x,v}(0) = x$, $\dot{\sigma}_{x,v}(0) = v$. If the manifold is endowed with a Riemannian structure and it holds that $\ddot{\sigma}_{x,v}(0) = 0$ for any x and v , then the retraction is said to be of second order.

2.2 DLRA

To explain the basic idea of DLRA, let us consider an initial value problem governed by a smooth vector field $F : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^{m \times n}$:

$$\begin{cases} A'(t) = F(A(t)), & t \in [0, T], \\ A(0) = A_0 \in \mathbb{R}^{m \times n}. \end{cases} \tag{2.2}$$

DLRA aims at finding an approximation of the so-called *ambient solution* $A(t) \in \mathbb{R}^{m \times n}$ on the manifold of rank- r matrices, for fixed $r \ll \min\{m, n\}$, in order to improve computational efficiency while maintaining satisfactory accuracy. Indeed, representing the approximation of $A(t)$ in factored form drastically reduces storage complexity. The central challenge of DLRA consists in computing efficiently a factored low-rank approximation of the ambient solution without having to first estimate the ambient solution and then to truncate it to an accurate rank- r approximation. In the following, the rank is fixed a priori and does not change throughout the integration interval. Then, the DLRA problem under consideration associated with the ambient Eq. (2.2) can be formalized as follows.

Problem 2.1 Given a smooth vector field $F : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^{m \times n}$, an initial matrix $A_0 \in \mathbb{R}^{m \times n}$ and a target rank r such that $\sigma_r(A_0) > \sigma_{r+1}(A_0)$, the DLRA problem consists in determining $t \mapsto Y(t) \in \mathcal{M}_r$ solving the following initial value problem

$$\begin{cases} \dot{Y}(t) = \Pi(Y(t))F(Y(t)), & t \in [0, T], \\ Y(0) = Y_0, \end{cases} \tag{2.3}$$

where $Y_0 = \Pi_{\mathcal{M}_r}(A_0) \in \mathcal{M}_r$, the rank- r truncated singular value decomposition of A_0 .

The origins of the above problem are rooted in the Dirac–Frenkel variational principle, by which the dynamics of (2.2) are optimally projected onto the tangent space of the manifold. In fact, for any $Y \in \mathcal{M}_r$, the orthogonal projection of $F(Y)$ to $T_Y\mathcal{M}_r$ returns the closest vector in the tangent space:

$$\|F(Y) - \Pi(Y)F(Y)\|_F = \min_{V \in T_Y\mathcal{M}_r} \|F(Y) - V\|_F. \tag{2.4}$$

This optimality criterion together with the optimal choice of Y_0 are the local first-order approximation of the computationally demanding optimality $Y_{\text{opt}}(t) = \Pi_{\mathcal{M}_r}(A(t))$. The gap (2.4) between the original dynamics and the projected dynamics of Problem 2.1 is known as the *modeling error* [19]. Given appropriate smoothness requirements on F and assuming the modeling error can be uniformly bounded in a neighborhood $\mathcal{U} \subset \mathbb{R}^{m \times n}$ of the trajectory $Y([0, T]) \subset \mathcal{M}$ of the exact solution of (2.3) as

$$\max_{Y \in \mathcal{U} \cap \mathcal{M}_r} \|F(Y) - \Pi(Y)F(Y)\|_F \leq \varepsilon, \tag{2.5}$$

then it can be shown, see e.g. [19, Theorem 2], that there exists a constant $C > 0$ depending on the final time T such that

$$\|A(T) - Y(T)\|_F \leq C(\delta_0 + \varepsilon), \tag{2.6}$$

where $\delta_0 = \|A(0) - Y(0)\|_F$ denotes the initial approximation error.

In recent years, several computationally efficient numerical integration schemes to approximate the solution to Problem 2.1 have been proposed [8, 19, 24]. Their output is a time discretization of the solution, where $Y_k \in \mathcal{M}_r$ approximates $Y(k\Delta t)$, for every $k = 0, \dots, N$, assuming—for simplicity—that a fixed step size $\Delta t = T/N$ is used. Existing error analysis results [18, 19] state that, provided the step size is chosen small enough, the error at the final time can be bounded as

$$\|Y_N - A(T)\|_F \leq \tilde{C}(\delta_0 + \varepsilon + \Delta t^q), \tag{2.7}$$

for some integer $q \geq 1$ and a constant $\tilde{C} > 0$ that depends on the final time and the problem at hand. The constant q is called the *convergence order* of the time stepping method.

The most direct strategy to numerically integrate Problem 2.1 is to write Y in rank- r factorization and derive individual evolution equations for the factors [20]. However, the computational advantages of such an approach are undermined by the high stiffness of the resulting equations when the r th singular value of the approximation becomes small, which is often the case in applications. In turn, this enforces unreasonably small step sizes in order to guarantee the stability of explicit integrators. The projector-splitting schemes proposed by Lubich and Oseledets [24] were the first to remedy this issue. Since then, a collection of methods have been designed that achieve stability independently of the presence of small singular values [7–9, 17, 18]. As we show in the following section, some of these DLRA algorithms are directly related to retractions.

3 DLRA algorithms and low-rank retractions

The concept of retraction was initially formulated in the context of numerical time integration on manifolds. In the work of Shub [27], what is now called a retraction was proposed as a generic tool to develop extensions of the forward Euler method that incorporate manifold constraints. Indeed, any retraction admits the following first-order approximation property. Let $t \mapsto Y(t)$ denote the exact solution to a manifold-constrained ODE such as (2.3) but on a general embedded manifold \mathcal{M} . Provided the vector field governing the ODE is sufficiently smooth, the defining properties of a retraction imply that [27, Theorem 1]

$$\|Y(t + \tau) - R_{Y(t)}(\tau \dot{Y}(t))\| = O(\tau^2). \quad (3.1)$$

This mirrors the local truncation error achieved by the forward Euler method on a Euclidean space while ensuring that the curve $\tau \mapsto R_{Y(t)}(\tau \dot{Y}(t))$ remains on the constraint manifold. Hence, if Y_k indicates an approximation to the exact solution $Y(t_k)$ at time $t = t_k$, the update rule

$$Y_{k+1} = R_{Y_k}(\Delta t \dot{Y}_k) \quad (3.2)$$

for approximating the solution at $t = t_{k+1} = t_k + \Delta t$, for a given step size $\Delta t > 0$, is a retraction-based generalization of the forward Euler for manifold-constrained ODEs. By virtue of (3.1), for any retraction on \mathcal{M} , the scheme (3.2) achieves local truncation error of order two (see Sect. 4 below for the definition of the order), which implies global first-order convergence under certain conditions [13, Section II.3].

Let us now specialize the discussion on the relation between retractions and numerical integration of manifold-constrained ODEs to the DLRA differential equation of Problem 2.1 evolving on the fixed-rank manifold \mathcal{M}_r . Implementing the general scheme (3.2) requires the choice of a particular low-rank retraction. A substantial number of retractions are known for the fixed-rank manifold [4] and, as discussed in the following, several existing DLRA integration schemes are realized as (3.2) for a particular choice of retraction on \mathcal{M}_r .

3.1 Metric-projection retraction

A large class of methods to numerically integrate manifold-constrained ODEs fit into the class of *projection methods*; see [12, §IV.4]. Each integration step is carried out in the embedding space with a Euclidean time-stepping method and is followed by the metric projection onto the constraint manifold. For the DLRA differential equation of Problem 2.1, projection methods have been studied in [19]. One step of the *projected forward Euler method* takes the form

$$Y_{k+1} = \Pi_{\mathcal{M}_r}(Y_k + \Delta t \Pi(Y_k)F(Y_k)), \tag{3.3}$$

where $\Pi_{\mathcal{M}_r}$ coincides with the rank- r truncated SVD. This integration scheme fits into the general class of retraction-based schemes (3.2). Indeed, given $X \in \mathcal{M}_r$ and $Z \in T_X \mathcal{M}_r$, the metric projection of the ambient space point $X + Z$ defines a retraction [3, §3],

$$R_X^{\text{SVD}}(Z) := \Pi_{\mathcal{M}_r}(X + Z), \tag{3.4}$$

which we refer to as the SVD retraction. Since $Z \in T_X \mathcal{M}_r$, the matrix $X + Z$ is of rank at most $2r$. This allows for an efficient implementation of this retraction, as detailed in Algorithm 1.

The projected forward Euler method has order $q = 1$ [19, Theorem 4]. To achieve higher orders of convergence, *projected Runge–Kutta* (PRK) methods have been proposed [19, §5]. The intermediate stages of such methods are obtained by replacing the forward Euler update in (3.3) with the updates of an explicit Runge–Kutta method. Although the update vectors may not belong to a single tangent space, the PRK recursion can still be written because the SVD retraction has the particular property that it remains well-defined not only for tangent vectors but also for sufficiently small general vectors $Z \in \mathbb{R}^{m \times n}$ (a property that makes it an *extended retraction* [4, §2.3]). We shall abbreviate the PRK method with $s \geq 1$ stages as PRK s . Note that PRK1 coincides with the projected forward Euler method above.

Algorithm 1 SVD retraction on \mathcal{M}_r .

Input: $X = U_0 \Sigma_0 V_0^\top \in \mathcal{M}_r$, $Z = U M V^\top + U_p V_p^\top + U V_p^\top \in T_X \mathcal{M}_r$.

- 1: $Q_U R_U = U_p$ with Q_U orthonormal;
 - 2: $Q_V R_V = V_p$ with Q_V orthonormal;
 - 3: $[U_S, \Sigma_S, V_S] = \text{SVD} \left(\begin{bmatrix} \Sigma + M & R_V^\top \\ R_U & 0 \end{bmatrix} \right)$;
 - 4: $\Sigma_1 = \Sigma_S(1:r, 1:r)$;
 - 5: $U_1 = [U_0 \ Q_U] U_S(:, 1:r)$;
 - 6: $V_1 = [V_0 \ Q_V] V_S(:, 1:r)$;
 - 7: **return** : $U_1 \Sigma_1 V_1^\top =: R_X^{\text{SVD}}(Z)$;
-

3.2 Projector-splitting KSL retraction

The expression (2.1) for the orthogonal projection of $Z \in \mathbb{R}^{m \times n}$ onto the tangent space at $Y = U \Sigma V^T \in \mathcal{M}_X$ can be expressed as

$$\Pi(Y)Z = ZVV^T - UU^T ZVV^T + UU^T Z. \tag{3.5}$$

The *projector-splitting* scheme for DLRA introduced in [24] is derived by applying this decomposition to the right-hand side of (2.3) and using standard Lie–Trotter splitting. Each integration step comprises three integration substeps identified with the letters K, S and L, corresponding respectively to the three terms in (3.5). When each of the integration substeps is performed with a forward Euler update, we refer to this scheme as the KSL scheme. The scheme is proved to be first-order accurate independently of the presence of small singular values [18, Theorem 2.1].

As shown in [4, Theorem 3.3], one step of the KSL scheme actually defines a second-order retraction for the fixed-rank manifold. In fact, it coincides modulo third-order terms with the orthographic retraction presented in Sect. 3.3.1. The KSL retraction is denoted by R^{KSL} and its computation is summarized in Algorithm 3. Hence, the KSL integration scheme fits into the general scheme (3.2) for Problem 2.1 as it can simply be written as

$$Y_{k+1} = R_{Y_k}^{\text{KSL}}(\Delta t \Pi(Y_k) F(Y_k)). \tag{3.6}$$

Algorithm 2 KLS retraction

Input: $X = U_0 \Sigma_0 V_0^T \in \mathcal{M}_r$, $Z = U M V^T + U_p V^T + U V_p^T \in T_X \mathcal{M}_r$.

- 1: (K-step) $U_1 S_U = U_0(\Sigma_0 + M) + U_p$ with U_1 orthonormal;
 - 2: (L-step) $V_1 S_V = V_0(\Sigma_0^T + M^T) + V_p$ with V_1 orthonormal;
 - 3: $L = U_1^T U_0$;
 - 4: $R = V_1^T V_0$;
 - 5: (S-step) $\Sigma_1 = L \left[(\Sigma_0 + M) R^T + V_p^T V_1 \right] + U_1^T U_p R^T$; ▷ equivalent to $U_1^T (X + Z) V_1$
 - 6: Optional: $[\tilde{U}_1, \tilde{\Sigma}_1, \tilde{V}_1] = \text{SVD}(\Sigma_1)$;
 - 7: Optional: $U_1 \leftarrow \tilde{U}_1, \Sigma_1 \leftarrow \tilde{\Sigma}_1, V_1 \leftarrow \tilde{V}_1$;
 - 8: **return** : $U_1 \Sigma_1 V_1^T =: R_X^{\text{KLS}}(Z)$;
-

3.3 Projector-splitting KLS retraction

Recently, a modification to KSL projector-splitting method [8] was proposed to improve its (parallel) performance, while maintaining the stability and accuracy properties of the KSL method. The scheme is known as the *unconventional integrator* and it is a modification of the KSL scheme where the L-step is performed before the S-step. Hence, when each of the integration substeps is performed with a forward Euler update, we refer to it as the KLS scheme. The KLS scheme comes with the computational advantage of allowing the K-step and L-step to be performed in parallel without

compromising first-order accuracy and stability with respect to small singular values [8, Theorem 4].

As we prove in the next section, one step of the KLS scheme also defines a retraction for the fixed-rank manifold. We denote it by R^{KLS} and its computation is detailed in Algorithm 2. Then, as for other schemes so far, the KLS integration scheme for DLRA takes the simple form

$$Y_{k+1} = R_{Y_k}^{KLS} (\Delta t \Pi(Y_k) F(Y_k)). \tag{3.7}$$

Algorithm 3 KSL retraction

Input: $X = U_0 \Sigma_0 V_0^T \in \mathcal{M}_r, Z = U M V^T + U_p V^T + U V_p^T \in T_X \mathcal{M}_r.$

- 1: (K-step) $U_1 \tilde{\Sigma}_1 = U_0(\Sigma_0 + M) + U_p$ with U_1 orthonormal;
 - 2: (S-step) $\tilde{\Sigma}_0 = \tilde{\Sigma}_1 - (U_1^T U_p + (U_1^T U_0) M);$
 - 3: (L-step) $V_1 \Sigma_1^T = V_0 \tilde{\Sigma}_0^T + Z^T U_1$ with V_1 orthonormal;
 - 4: Optional: $[\tilde{U}_1, \tilde{\Sigma}_1, \tilde{V}_1] = \text{SVD}(\tilde{\Sigma}_1);$
 - 5: Optional: $U_1 \leftarrow U_1 \tilde{U}_1, \Sigma_1 \leftarrow \tilde{\Sigma}_1, V_1 \leftarrow V_1 \tilde{V}_1;$
 - 6: **return** : $U_1 \Sigma_1 V_1^T =: R_X^{KSL}(Z);$
-

3.3.1 The KLS retraction and the orthographic retraction

The goal of this section is to prove that the update rule of the KLS scheme, as detailed by Algorithm 2, defines a second-order retraction. The strategy to reach this goal is based on the observation that the KLS update is obtained as a small perturbation of another commonly encountered retraction, known as the orthographic retraction.

Algorithm 4 Orthographic retraction

Input: $X = U_0 \Sigma_0 V_0^T \in \mathcal{M}_r, Z = U M V^T + U_p V^T + U V_p^T \in T_X \mathcal{M}_r.$

- 1: $U_1 S_U = U_0(\Sigma_0 + M) + U_p$ with U_1 orthonormal;
 - 2: $V_1 S_V = V_0(\Sigma_0^T + M^T) + V_p$ with V_1 orthonormal;
 - 3: $\Sigma_1 = S_U(\Sigma_0 + M)^{-1} S_V^T,$
 - 4: Optional: $[\tilde{U}_1, \tilde{\Sigma}_1, \tilde{V}_1] = \text{SVD}(\Sigma_1);$
 - 5: Optional: $U_1 \leftarrow U_1 \tilde{U}_1, \Sigma_1 \leftarrow \tilde{\Sigma}_1, V_1 \leftarrow V_1 \tilde{V}_1;$
 - 6: **return** : $U_1 \Sigma_1 V_1^T =: R_X^{\text{ORTH}}(Z);$
-

The orthographic retraction. Specialized to the fixed-rank matrix manifold \mathcal{M}_r , the orthographic retraction consists in perturbing the point $X \in \mathcal{M}_r$ in the ambient space as $X + Z \in \mathbb{R}^{m \times n}$ and projecting back onto the manifold but only along vectors from the normal space of the starting point, see Fig. 1. Formally this reads:

$$R_X^{\text{ORTH}}(Z) = \arg \min_{Y \in (X + Z + N_X \mathcal{M}_r) \cap \mathcal{M}_r} \|X + Z - Y\|_F. \tag{3.8}$$

A closed-form expression for the solution of this optimization problem for the case of \mathcal{M}_r is established in [4, §3.2] and is the output of Algorithm 4. By virtue of the

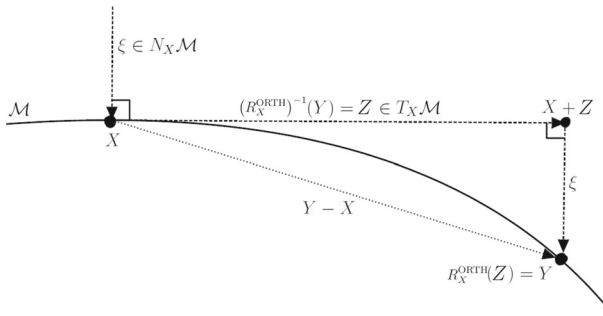


Fig. 1 Orthographic retraction and its inverse

analysis carried out in [3], the orthographic retraction is a second-order retraction. A remarkable property of the orthographic retraction is that its local inverse can be computed easily. As suggested by the construction illustrated in Fig. 1, the inverse orthographic retraction is obtained by projecting the ambient space difference onto the tangent space, i.e.,

$$\left(R_X^{\text{ORTH}}\right)^{-1}(Y) = \Pi(X)(Y - X). \tag{3.9}$$

This opens the possibility to use the inverse orthographic retraction in practical numerical procedures such as Hermite interpolation of a manifold curve [26] used in the definition of the Projected Ralston–Hermite scheme in Sect. 4.2.

A careful inspection of the computations for the orthographic retraction reported in Algorithm 4 reveals that it is identical to the update rule of the KLS scheme given in Algorithm 2, up to the computation of Σ_1 . Let Σ_1^{KLS} and Σ_1^{ORTH} denote the quantities computed respectively at step 5 of Algorithm 2 and step 3 of Algorithm 4. Note that Σ_1^{ORTH} can be rewritten without explicitly computing the factors S_U and S_V . Using the relations

$$S_U = L(\Sigma_0 + M) + U_1^\top U_p, \quad S_V = R(\Sigma_0^\top + M^\top) + V_1^\top V_p, \tag{3.10}$$

with $L = U_1^\top U_0$, $R = V_1^\top V_0$ we evince that

$$\begin{aligned} \Sigma_1^{\text{ORTH}} &= S_U(\Sigma_0 + M)^{-1} S_V^\top \\ &= L \left[(\Sigma_0 + M) R^\top + V_p^\top V_1 \right] + U_1^\top U_p R^\top + U_1^\top U_p (\Sigma_0 + M)^{-1} V_p^\top V_1. \end{aligned} \tag{3.11}$$

This highlights that the quantity Σ_1 computed in the KLS scheme and the orthographic retraction differ by

$$\Sigma_1^{\text{ORTH}} - \Sigma_1^{\text{KLS}} = U_1^\top U_p (\Sigma_0 + M)^{-1} V_p^\top V_1. \tag{3.12}$$

Using this observation together with the second-order property of the orthographic retraction allows us to show that the KLS procedure defines a second-order retraction. This link with the orthographic retraction mirrors the same observation made for the closely related KSL scheme [4, Theorem 3.3].

Proposition 3.1 *The procedure of Algorithm 2 defines a second-order retraction (called the KLS retraction).*

Proof The proof relies on a necessary and sufficient condition for a retraction to be a second-order retraction stated in [3, Proposition 3]. On a general Riemannian manifold \mathcal{M} , a mapping $R : T\mathcal{M} \rightarrow \mathcal{M}$ is a second-order retraction if and only if for all $(x, v) \in T\mathcal{M}$ the curve $t \mapsto R_x(tv)$ is well-defined for all sufficiently small t and satisfies

$$R_x(tv) = \text{Exp}_x(tv) + O(t^3), \tag{3.13}$$

where $t \mapsto \text{Exp}_x(tv)$ is a parametrization of the geodesic passing through x with velocity v obtained with the Riemannian exponential map at x [5, Definition 10.16].

Given $X \in \mathcal{M}_r$ and $Z \in T_X\mathcal{M}_r$, the orthonormalizations of the first two lines of the KLS scheme in Algorithm 2 are uniquely defined provided the matrices to orthonormalize have full rank. This is the case for any Z sufficiently small by lower semi-continuity of the matrix rank. By smoothness of the orthonormalization process, we know that the matrix Σ_1 depends smoothly on Z . Since for $Z = 0$ we have $\Sigma_1 = \Sigma_0$, assumed to be full rank, the matrix Σ_1 has full rank for sufficiently small Z . Hence $U_1 \Sigma_1 V_1^\top = R_X^{\text{KLS}}(Z)$ is uniquely and smoothly defined for any Z in a neighborhood of the origin of $T_X\mathcal{M}_r$ and belongs to \mathcal{M}_r .

Now, consider the curves $t \mapsto R_X^{\text{ORTH}}(tZ)$ and $t \mapsto R_X^{\text{KLS}}(tZ)$, well-defined for sufficiently small t . Since these curves share the left and right singular vectors $U_1(t)$ and $V_1(t)$, their difference is given by

$$R_X^{\text{ORTH}}(tZ) - R_X^{\text{KLS}}(tZ) = t^2 U_1(t)^\top U_p(\Sigma_0 + tM)^{-1} V_p^\top V_1(t) = t^2 C(t),$$

where $C(t) := U_1(t)^\top U_p(\Sigma_0 + tM)^{-1} V_p^\top V_1(t)$. Let us show that $C(t) = o(t)$, i.e. $\lim_{t \rightarrow 0} C(t)/t = 0$. Since $U_1(0) = U_0$ and $V_1(0) = V_0$, by definition of tangent space of \mathcal{M}_r , we know that $U_1(0)^\top U_p = 0$ and $V_p^\top V_1(0) = 0$. Hence $C(0) = 0$ and therefore

$$\lim_{t \rightarrow 0} \frac{C(t)}{t} = \lim_{t \rightarrow 0} \frac{C(t) - C(0)}{t}. \tag{3.14}$$

This coincides with $C'(0)$ since C is smooth for small t . But since

$$C'(0) = \underbrace{U_1'(0)^\top U_p \Sigma_0^{-1} V_p^\top V_1(0)}_{=0} + \underbrace{U_1(0)^\top U_p}_{=0} \frac{d}{dt} (\Sigma_0 + tM)^{-1} \Big|_{t=0} \underbrace{V_p^\top V_1(0)}_{=0} \tag{3.15}$$

$$+ \underbrace{U_1(0)^\top U_p \Sigma_0^{-1} V_p^\top V_1'(0)}_{=0} = 0 \tag{3.16}$$

we can infer that $C(t) = o(t)$ and therefore that $R_X^{KLS}(tZ) = R_X^{ORTH}(tZ) + o(t^3)$. Using the result [3, Proposition 2.3] explained in the beginning of the proof and the second-order property of the orthographic retraction, we obtain that the retraction curve $t \mapsto R_X^{ORTH}(tZ)$ approximates the geodesic $t \mapsto \text{Exp}_X(tZ)$, as $R_X^{ORTH}(tZ) = \text{Exp}_X(tZ) + O(t^3)$. Combining the two results leads to

$$R_X^{KLS}(tZ) = R_X^{ORTH}(tZ) + o(t^3) = \gamma_{X,Z}(t) + O(t^3) + o(t^3) = \gamma_{X,Z}(t) + O(t^3). \tag{3.17}$$

Using once again the result from [3, Proposition 2.3], this implies that R_X^{KLS} is indeed a second-order retraction. □ □

4 New retraction-based time stepping algorithms

Having seen in Sect. 3 that existing DLRA time integration algorithms can be directly related with particular low-rank retractions, we will now discuss the opposite direction, how low-rank retractions can be used to produce new integration schemes.

Before proceeding to the derivation of the new methods, we first briefly describe the rationale behind these methods for the initial value problem (2.2) evolving in $\mathbb{R}^{m \times n}$. We let $A(\Delta t)$ denote the exact solution at time $t = \Delta t$ and $\tilde{A}(\Delta t)$ its approximation obtained by performing one step of a given numerical integration scheme starting from $A(0)$. The scheme is said to have order q if the local truncation error satisfies

$$\|A(\Delta t) - \tilde{A}(\Delta t)\| = O(\Delta t^{q+1}). \tag{4.1}$$

In other words, any curve $\tilde{A}(\cdot)$ on $\mathbb{R}^{m \times n}$ that is a sufficiently good approximation of $A(\cdot)$ in the vicinity of 0 defines a suitable numerical integration scheme. In the two methods we propose, the curves \tilde{A} are constructed through manifold interpolation, based on retractions.

By definition, see Sect. 2.1, any retraction induces manifold curves with prescribed initial position and velocity. If, additionally, the retraction has second order, it is possible to prescribe an initial acceleration.

Proposition 4.1 ([5, Exercise 5.46], [29, Corollary 3.2]) *Let R be a retraction on \mathcal{M} and consider arbitrary $x \in \mathcal{M}$ and $v, a \in T_x \mathcal{M}$. Then the curve*

$$\sigma(t) = R_x \left(tv + \frac{t^2}{2} a \right) \tag{4.2}$$

is well-defined for all t sufficiently small and satisfies

1. $\sigma(0) = x$ and $\dot{\sigma}(0) = v$,
2. $\ddot{\sigma}(0) = a$ if R is a second-order retraction.

When setting $x = \gamma(0)$, $v = \dot{\gamma}(0)$ and $a = \ddot{\gamma}(0)$ for a smooth manifold curve γ , it is shown in [29, Proposition 3.22] that the Riemannian distance between γ and the retraction curve constructed by Proposition 4.1 satisfies

$$d(\gamma(t), \sigma(t)) = O(t^3), \quad t \rightarrow 0, \tag{4.3}$$

if R is a second-order retraction.

4.1 An accelerated forward Euler scheme

The first method we propose for DLRA is obtained by accelerating the projected forward Euler scheme (3.3) through a second-order correction.

4.1.1 Euclidean analog

For illustration, let us first derive the accelerated forward Euler method for the initial value problem (2.2) evolving in $\mathbb{R}^{m \times n}$. If the solution to (2.2) is sufficiently smooth, we get from its Taylor expansion that

$$A(\Delta t) - \left(A(0) + \Delta t A'(0) + \frac{\Delta t^2}{2} A''(0) \right) = O(\Delta t^3).$$

Assuming we can compute $A''(0) = DF(A_0)[F(A_0)]$, the differential of the vector field along $F(A_0)$, the Euclidean accelerated forward Euler (AFE) scheme is defined by the update

$$A^{\text{AFE}}(\Delta t) = A_0 + \Delta t F(A_0) + \frac{\Delta t^2}{2} DF(A_0)[F(A_0)]. \tag{4.4}$$

The Euclidean AFE scheme is of order $q = 2$ since the local truncation error is by construction of order $O(\Delta t^3)$. The scheme is conditionally stable with the same stability domain as any 2-stages explicit Runge–Kutta method of order 2 [14, Theorem 2.2]. That being said, this scheme requires the evaluation of the Jacobian at each step, which renders the AFE scheme non-competitive on Euclidean spaces. In fact, for Euclidean ODEs there are simple second-order methods using only the evaluation of the vector field, such as Runge–Kutta methods.

In the following, we generalize the AFE scheme for the DLRA differential equation of Problem 2.1 by first showing how to compute the acceleration of the exact solution. As we will see below, in Proposition 4.3, it is the sum of two terms: the tangent component of the ambient acceleration and the acceleration due to the curvature of the manifold. The latter is expressed using the *Weingarten map*, a classical concept of Riemannian geometry that we briefly introduce in the next section.

4.1.2 The Weingarten map

As before, we consider a Riemannian submanifold \mathcal{M} embedded into a finite-dimensional Euclidean space \mathcal{E} ; see [23, §8] for a more general setting. We let $\mathfrak{X}_T(\mathcal{M})$ and $\mathfrak{X}_N(\mathcal{M})$ denote the set of smooth tangent and normal vector fields on \mathcal{M} , respectively.

The Weingarten map is constructed from the second fundamental form. The latter measures the discrepancy between the ambient Euclidean connection and the Riemannian connection ∇ on the submanifold \mathcal{M} . For every $W, Z \in \mathfrak{X}_T(\mathcal{M})$ and $x \in \mathcal{M}$, we have that $\nabla_W Z(x) = \Pi(x)D_W Z(x)$, where D is the Euclidean connection corresponding to standard directional derivative. Hence, the vector

$$D_W Z(x) - \nabla_W Z(x) = D_W Z(x) - \Pi(x)D_W Z(x) \tag{4.5}$$

belongs to the normal space at x and depends smoothly on x . Hence, the function $(I - \Pi)D_W Z$ is a smooth normal vector field on \mathcal{M} .

Definition 4.1 The second fundamental form $\mathbb{I} : \mathfrak{X}_T(\mathcal{M}) \times \mathfrak{X}_T(\mathcal{M}) \rightarrow \mathfrak{X}_N(\mathcal{M})$ is a symmetric bilinear map defined by $\mathbb{I}(W, Z) = (I - \Pi)D_W Z$.

Definition 4.2 The Weingarten map is

$$\begin{aligned} \mathcal{W} : \mathfrak{X}_T(\mathcal{M}) \times \mathfrak{X}_N(\mathcal{M}) &\rightarrow \mathfrak{X}_T(\mathcal{M}) \\ (W, N) &\mapsto \mathcal{W}(W, N) \end{aligned} \tag{4.6}$$

defined by the multilinear form $\langle N, \mathbb{I}(W, Z) \rangle = \langle \mathcal{W}(W, N), Z \rangle, \forall Z \in \mathfrak{X}_T(\mathcal{M})$.

Since $\nabla_W Z$ can be shown to depend only on the value of the vector field W at x , using the properties of the Levi-Civita connection both the second fundamental form and the Weingarten map can be defined pointwise, i.e. depending only on values of the vector fields at x [23, Proposition 8.1]. Given $w, z \in T_x \mathcal{M}$ and $n \in N_x \mathcal{M}$, the Weingarten map $\mathcal{W}_x : T_x \mathcal{M} \times N_x \mathcal{M} \rightarrow T_x \mathcal{M}$ is defined pointwise as

$$\langle \mathcal{W}_x(w, n), z \rangle_x = \langle \mathcal{W}(W, N), Z \rangle \Big|_x \tag{4.7}$$

for any $W, Z \in \mathfrak{X}_T(\mathcal{M})$ and $N \in \mathfrak{X}_N(\mathcal{M})$ satisfying $W(x) = w, Z(x) = z$ and $N(x) = n$. The following characterization relates the Weingarten map at x with the differential of the tangent space projection.

Proposition 4.2 ([2, Theorem 1]) *For any $x \in \mathcal{M}, w \in T_x \mathcal{M}$ and $n \in N_x \mathcal{M}$, the Weingarten map satisfies*

$$\mathcal{W}_x(w, n) = D\Pi(x)[w]n = \Pi(x)D\Pi(x)[w]n = \Pi(x)D\Pi(x)[w] \Pi^\perp(x)q \tag{4.8}$$

for any $q \in T_x \mathcal{E} \simeq \mathcal{E}$ such that $\Pi^\perp(x)q = n$.

Expression for the fixed-rank manifold. Depending on the conventions used to represent points and tangent vectors, many equivalent expressions are known for the Weingarten map of the fixed-rank manifold, see for example [2, 11]. In our conventions, for any $Y = U\Sigma V^T \in \mathcal{M}_r$, $T = U M V^T + U_p V^T + U V_p^T \in T_Y \mathcal{M}_r$ and $N \in N_Y \mathcal{M}_r$ the Weingarten map can be computed as

$$\mathcal{W}_Y(T, N) = N V_p \Sigma^{-1} V^T + U \Sigma^{-1} U_p^T N \in T_Y \mathcal{M}_r. \tag{4.9}$$

This expression nicely highlights two notable features that are known for the fixed-rank manifold: it is a ruled surface with unbounded curvature. In fact, along the subspace associated to the $U M V^T$ term, \mathcal{M}_r is flat, while the curvature along the other directions grows unbounded as $\sigma_r(Y) \rightarrow 0$.

4.1.3 The accelerated forward Euler (AFE) integration scheme for DLRA

With the definitions introduced in Sect. 4.1.2, we can compute the acceleration of the DLRA solution curve allowing to generalize the AFE integration scheme to Problem 2.1.

Proposition 4.3 *If a smooth curve on \mathcal{M}_r is defined by $\dot{Y} = \Pi(Y)F(Y)$ for some smooth vector field $F : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^{m \times n}$, then its intrinsic acceleration can be computed as*

$$\ddot{Y} = \Pi(Y)DF(Y)[\Pi(Y)F(Y)] + \mathcal{W}_Y(\Pi(Y)F(Y), \Pi^\perp(Y)F(Y)). \tag{4.10}$$

Proof By definition of the Levi-Civita connection for Riemannian submanifolds we have

$$\ddot{Y} = \nabla_{\dot{Y}} \dot{Y} = \Pi(Y)D_{\dot{Y}}(\Pi(Y)F(Y)). \tag{4.11}$$

Using the product rule and the fact that $\Pi(Y)^2 = \Pi(Y)$, it follows from the last equality of Proposition 4.2 that

$$\begin{aligned} & \Pi(Y)D_{\dot{Y}}(\Pi(Y)F(Y)) \\ &= \Pi(Y) \left(\Pi(Y)DF(Y)[\dot{Y}] + \Pi(Y)D\Pi(Y)[\dot{Y}]F(Y) \right) \\ &= \Pi(Y)DF(Y)[\Pi(Y)F(Y)] + \Pi(Y)D\Pi(Y)[\Pi(Y)F(Y)]\Pi^\perp(Y)F(Y) \\ &= \Pi(Y)DF(Y)[\Pi(Y)F(Y)] + \mathcal{W}_Y(\Pi(Y)F(Y), \Pi^\perp(Y)F(Y)). \end{aligned}$$

□

Hence, to mimic the Euclidean AFE update defined by Eq. (4.4), we need to construct a smooth manifold curve $Y^{\text{AFE}}(\Delta t)$ such that

$$Y^{\text{AFE}}(0) = Y_0, \quad \dot{Y}^{\text{AFE}}(0) = \dot{Y}(0), \quad \ddot{Y}^{\text{AFE}}(0) = \ddot{Y}(0).$$

Proposition 4.1 shows that we can construct such a curve using a second-order retraction. Let $\mathbf{R}^{\mathbb{I}}$ indicate any second-order retraction, then the manifold analogues of the AFE update (4.4) reads as

$$Y^{\text{AFE}}(\Delta t) = \mathbf{R}_{Y_0}^{\mathbb{I}} \left(\Delta t \dot{Y}_0 + \Delta t^2 \ddot{Y}_0 / 2 \right). \quad (4.12)$$

Then, the AFE scheme for DLRA takes the form

$$Y_{k+1} = \mathbf{R}_{Y_k}^{\mathbb{I}} \left(\Delta t \Pi(Y_k) F(Y_k) + \frac{\Delta t^2}{2} \ddot{Y}_k \right) \quad (4.13)$$

with

$$\ddot{Y}_k = \Pi(Y_k) D F(Y_k) [\Pi(Y_k) F(Y_k)] + \mathcal{W}_{Y_k} \left(\Pi(Y_k) F(Y_k), \Pi^\perp(Y_k) F(Y_k) \right). \quad (4.14)$$

As a consequence of (4.3), this scheme admits a local truncation error of order $O(\Delta t^3)$.

All the retractions for the fixed-rank manifold presented in Sect. 3 have the second-order property. In principle, all of them are suited to be used as $\mathbf{R}^{\mathbb{I}}$ in (4.13), however, experiments reported in Sect. 5 suggest the orthographic retraction is the most convenient in terms of speed, accuracy and stability.

4.2 The projected Ralston–Hermite integration scheme

As put forth in [26], retractions can also be used to define a manifold curve with prescribed endpoints and endpoint velocities. For tangent bundle data points $(x_0, v_0), (x_1, v_1) \in T\mathcal{M}$ and some parameters $t_0 < t_1$, we denote by H the retraction-based Hermite (RH) interpolant defined by [26, Corollary 7], which satisfies

$$H(t_i) = x_i, \quad \dot{H}(t_i) = v_i, \quad i = 0, 1. \quad (4.15)$$

We will employ the notation $H(t; (p_0, v_0, t_0), (p_1, v_1, t_1))$ to underline the dependence of H on the interpolation data. The curve H is constructed by a manifold extension of the well-known De Casteljau algorithm. But instead of using geodesic segments as building blocks of the procedure, the RH interpolant is defined with endpoint curves constructed with the use of a retraction and its local inverse, making the method more efficient and more broadly applicable. An efficient evaluation of the curve H requires a retraction for which the inverse retraction admits a computationally affordable procedure to evaluate. On the fixed-rank manifold, the orthographic retraction is a suitable candidate to construct the RH interpolant.

As stated in [26, Theorem 17], the RH interpolant achieves $O(\Delta t^4)$ approximation error as $\Delta t \rightarrow 0$ in the case where the interpolation data is sampled from a smooth manifold curve at t_0 and $t_1 = t_0 + \Delta t$. The Projected Ralston–Hermite (PRH) scheme aims at leveraging this approximation power to define a numerical integration update

rule with small local truncation error. Let us derive the PRH method for the initial value problem (2.2) evolving in $\mathbb{R}^{m \times n}$.

4.2.1 Euclidean analog

Given $A_0, A_1, V_0, V_1 \in \mathbb{R}^{m \times n}$ and $t_0 < t_1$, the Hermite interpolant $\tau \mapsto H_P(\tau; (t_0, A_0, V_0), (t_1, A_1, V_1))$ is the unique third-degree polynomial curve in $\mathbb{R}^{m \times n}$ that satisfies $H_P(t_i) = A_i, H'_P(t_i) = V_i$, for $i = 0, 1$. The polynomial curve H_P can be used to define the following multistep integration scheme for the initial value problem (2.2):

$$A_{k+2} = H_P(t_{k+2}; (t_k, A_k, F(A_k)), (t_{k+1}, A_{k+1}, F(A_{k+1}))). \tag{4.16}$$

Working out the expression for H_P allows rewriting the recursive relation (4.16) as

$$A_{k+2} = 5A_k - 4A_{k+1} + \Delta t (2F(A_k) + 4F(A_{k+1})) \tag{4.17}$$

$$= A_{k+1} + \Delta t \left(2F(A_k) + 4F(A_{k+1}) - 5 \left(\frac{A_{k+1} - A_k}{\Delta t} \right) \right). \tag{4.18}$$

As pointed out in [13, §III.3], this scheme has a local truncation error $O(\Delta t^4)$, the highest possible order for explicit 2-step method using these terms. However, it is not zero-stable and thus does not produce a convergent scheme of order 3. Nevertheless, this update rule can be combined with suitably chosen intermediate steps to recover stability. Consider the family of multistep methods obtained by taking an intermediate forward Euler step of length $\alpha \in (0, 1)$ combined with the update rule (4.16) as

$$\begin{cases} A_{k+\alpha} = A_k + \alpha \Delta t F(A_k), \\ A_{k+1} = H_P(t_{k+1}; (t_k, A_k, F(A_k)), (t_{k+\alpha}, A_{k+\alpha}, F(A_{k+\alpha}))). \end{cases} \tag{4.19}$$

These schemes are all part of the family of explicit Runge–Kutta methods as it can be shown that

$$A_{k+1} = A_k + \Delta t \left[\left(1 + \frac{1}{\alpha} - \frac{1}{\alpha^2} \right) F(A_k) + \left(\frac{1}{\alpha^2} - \frac{1}{\alpha} \right) F(A_{k+\alpha}) \right]. \tag{4.20}$$

For any $\alpha \in (0, 1)$, the scheme satisfies the first-order conditions of Runge–Kutta methods. Choosing α to satisfy also the second-order conditions narrows down the family to the scheme with $\alpha = 2/3$. This scheme is an explicit second-order Runge–Kutta method known as the Ralston scheme and is defined by the following Butcher table.

$$\begin{array}{c|cc} 0 & & \\ \frac{2}{3} & \frac{2}{3} & \\ \hline \frac{1}{4} & \frac{3}{4} & \end{array}$$

4.2.2 The Projected Ralston–Hermite (PRH) integration scheme for DLRA

Expressed in the form (4.19), the update rule of the Ralston scheme consists of a forward Euler step followed by a Hermite interpolation step. We can leverage this formulation and the RH interpolant (4.15) to extend the Ralston method to the manifold setting. Let R denote any retraction and R^I denote any retraction that can be used in practice to evaluate the RH interpolant (i.e. whose inverse can be efficiently computed). To indicate which retraction is used to construct the retraction-based Hermite interpolant H , we add it to its list of arguments. The Projected Ralston–Hermite (PRH) scheme for Problem 2.1 is then defined as

$$\begin{cases} Y_{k+2/3} = R_{Y_k} \left(\frac{2}{3} \Pi(Y_k) F(Y_k) \right), \\ Y_{k+1} = H(t_{k+1}; (t_k, Y_k, \Pi(Y_k) F(Y_k)), (t_k + \frac{2}{3} \Delta t, Y_{k+2/3}, \\ \Pi(Y_{k+2/3}) F(Y_{k+2/3})), R^I). \end{cases} \quad (4.21)$$

A suitable candidate for both retractions is $R = R^I = R^{\text{ORTH}}$, the orthographic retraction. As experiments in Sect. 5 suggest, this generalization to the manifold setting of the Ralston scheme maintains its second-order accuracy.

4.2.3 The APRH integration scheme

For the sake of completeness, a third scheme can be obtained by combining the AFE and the PRH schemes. Replacing the intermediate forward Euler step of the PRH scheme with an AFE update defines what we call the Accelerated Projected Ralston–Hermite scheme (APRH). It is defined by the recursive relation

$$\begin{cases} Y_{k+2/3} = R_{Y_k}^{\Pi} \left(\frac{2}{3} \Delta t \Pi(Y_k) F(Y_k) + \frac{2\Delta t^2}{9} \ddot{Y}_k \right), \\ Y_{k+1} = H(t_{k+1}; (t_k, Y_k, \Pi(Y_k) F(Y_k)), (t_k + \frac{2}{3} \Delta t, Y_{k+2/3}, \\ \Pi(Y_{k+2/3}) F(Y_{k+2/3})), R^I), \end{cases} \quad (4.22)$$

where \ddot{Y}_k is given by (4.14).

5 Numerical experiments

In this section, we illustrate the performances of the accelerated forward Euler (AFE) method, the Projected Ralston–Hermite (PRH) method and the accelerated Projected Ralston–Hermite (APRH) method. Experiments were executed with Matlab 2022b on a laptop computer with Intel i7 CPU (1.8GHz with single-thread mode) with 8GB of RAM, 1MB of L2 cache and 8MB of L3 cache. The implementation leverages the differential geometry tools of the Manopt library [6]. In particular, the orthographic retraction provided by Manopt is used for the implementation of AFE, PRH and APRH. An implementation of the KSL and KLS retractions as described by Algorithms 3 and 2 were added to the fixed-rank manifold factory. For the implementation of the projected

Table 1 Summary of the acronyms of the methods considered in the numerical experiments

KSL	Projector-splitting integrator [4]
KLS	Unconventional integrator [8]
PRKs ($s = 1, 2, 3$)	Projected Runge–Kutta methods [19]
AFE	Accelerated forward Euler method, Sect. 4.1
PRH	Projected Ralston–Hermite method, Sect. 4.2
APRH	Accelerated Projected Ralston–Hermite method, Sect. 4.2.3

Runge–Kutta method of [19], we also added an implementation of the truncated SVD extended retraction, accepting as input tangent vectors of arbitrary tangent spaces. See Table 1 for a summary of the acronyms of the methods appearing in the numerical experiments. For the sake of completeness, we recall that an s -stage Projected Runge–Kutta (PRKs) method for given Butcher table, as provided in [19], is defined as follows

$$(\text{PRKs}) \begin{cases} Y_{k+1} = \mathcal{R}(Y_k + h \sum_{j=1}^s b_j \Pi(\mathcal{R}(Z_j))F(\mathcal{R}(Z_j))), \\ Z_j = Y_k + h \sum_{l=1}^{j-1} \Pi(\mathcal{R}(Z_l))F(\mathcal{R}(Z_l)), \end{cases}$$

where \mathcal{R} denotes the metric projection on the fixed-rank manifold \mathcal{M}_r . We refer to [19, §5] for a detailed description of its efficient implementation, along with a recapitulation of the Butcher table for each s -stage method up to the third order.

5.1 Differential Lyapunov equation

The modeling error (2.4) introduced by DLRA is associated with the normal component of the vector field of the original differential equation. The effect of this error on the performance of DLRA integrators can be assessed by considering a class of matrix differential equations, the so-called differential Lyapunov equations [30, §6.1], which take the form

$$\begin{cases} A'(t) = LA(t) + A(t)L^\top + Q, & t \in [0, T], \\ A(0) = A_0, \end{cases} \tag{5.1}$$

for some $A_0, L, Q \in \mathbb{R}^{n \times n}$. Indeed, if A_0 has rank exactly r and the matrix Q is zero, then $A(t)$ is also of rank- r for every $t \in [0, T]$ [15, Lemma 1.22]. Therefore, the norm of Q is proportional to the modeling error.

In the following, L is the discretized 1D Laplace operator on a uniform grid, that is, L is the tridiagonal matrix with -2 on the main diagonal and 1 on the first off-diagonals. For the source term, we take $Q = \eta \tilde{Q} / \|\tilde{Q}\|_F$ for some $\eta > 0$ and where \tilde{Q} is a full-rank matrix generated from its singular value decomposition with randomly chosen singular vectors and prescribed singular values, decaying as $\sigma_i(\tilde{Q}) = 10^{2-i}$, for $i = 1, \dots, n$. The initial condition is taken to be of rank exactly r and is also assembled from a randomly generated singular value decomposition with a prescribed

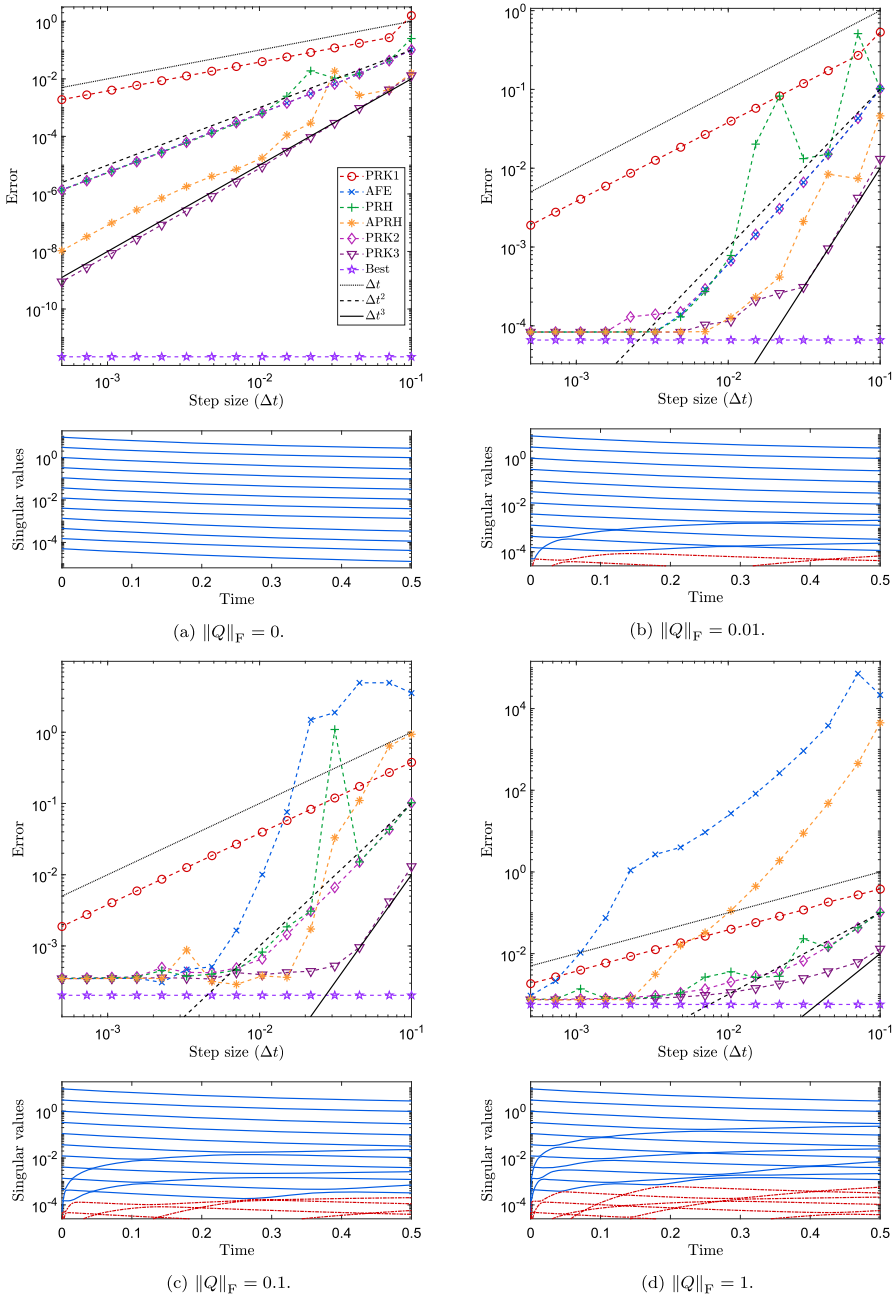


Fig. 2 Convergence of the error at final time for different DLRA integration schemes applied to the Lyapunov Eq. (5.1) with sources terms of different norms. The top plot in each panel is the final error $\|Y_{\Delta t}(T) - A(T)\|_2$ versus the step size Δt , where $Y_{\Delta t}$ is the approximation of A obtained with a step size Δt . The bottom plot reports the evolution of the singular values of the reference solution over time. The red dashed curves correspond to discarded singular values

geometric decay of non-zero singular values: $\sigma_i(A_0) = 3^{2-i}$, for all $i = 1, \dots, r$, and $\sigma_i(A_0) = 0$, for all $i = r + 1, \dots, n$.

In Fig. 2, we report the results with $n = 100$ and $r = 12$ of the following experiments. For different values of η , we numerically integrate the rank- r DLRA differential Eq. (2.3) applied to (5.1) with different numerical schemes and different time steps up to $T = 0.5$. A reference solution to the ambient Eq. (5.1) is found by using the MATLAB routine `ode45` between each time step, for a time step that is the smallest among those considered for the numerical integrators. We then plot as a function of the step size the 2-norm discrepancy between the reference solution at final time and its approximation obtained by numerical integration. The numerical results for the KSL and the KLS scheme were very similar to the ones of PRK1. Hence, they were omitted not to overcrowd the plots.

The panels of Fig. 2 correspond to the cases (a) $\eta = 0$, (b) $\eta = 0.01$, (c) $\eta = 0.1$, (d) $\eta = 1.0$. When the source term is zero, the reference solution is also of rank exactly r , as can be seen from the value of the best approximation error in panel (a). In this regime, the AFE and the PRH scheme both exhibit $O(\Delta t^2)$ error convergence, while the APRH scheme seem to reach an asymptotic $O(\Delta t^3)$ trend. The trade-off between accuracy and computational effort that can be seen in Fig. 3a shows that in this simple setting, the PRH, AFE and APRH schemes have comparable performances to PRK2. Turning on the source term determines a non-negligible best approximation error due to the growth of singular values that were initially zero, as can be seen in the bottom plots of panels (b), (c) and (d). The larger the source term's norm, the faster and the greater these singular values grow. Then, the numerical integrators converge to the exact solution of the projected system and so the error with respect to the ambient solutions stagnates at a value slightly higher than the best 2-norm approximation. While the PRH scheme preserves the $O(\Delta t^2)$ trend up to some oscillations as η increases, the AFE and APRH schemes seem to suffer instability when the normal component of the vector field is too large. In this more realistic scenario where the normal component of the vector field is non-negligible, only the PRH scheme remains comparable to PRK2 in terms of the trade-off between accuracy and effort, see Fig. 3b and Table 2.

5.2 Robustness to small singular values

A fundamental prerequisite for competitive DLRA integrators is to be resilient to the presence of small singular values in the solution. A detailed discussion on the topic can be found in [18]. In applications, very often the ambient solution admits an exponential decay of singular values. Hence, a good low-rank approximation is possible but the occurrence of small singular values is inevitable for DLRA to be accurate: a rank- r approximation of the solution must match the r th singular value of the ambient solution, which is small if the approximation error is small.

The smaller the singular values of the solution, the greater the stiffness of the DLRA differential Eq. (2.3): the Lipschitz constant of the vector field F gets multiplied by the Lipschitz constant of the tangent space projection, which is inversely proportional to the smallest non-zero singular value of the base point [20, Lemma 4.2]. Accordingly,

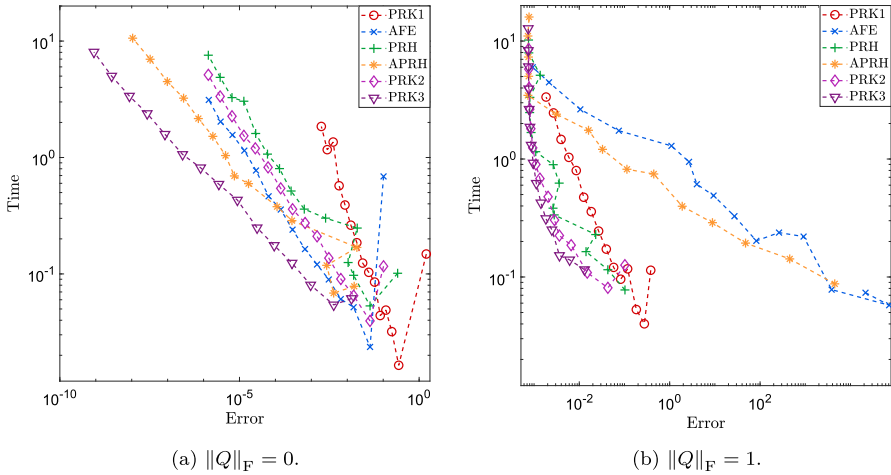


Fig. 3 Computational effort in terms of wall-clock time against the error with respect to the reference solution achieved by different numerical integration schemes and with different step sizes. The results were collected from the same experiment of panels (a) and (d) of Fig. 2

Table 2 Average time in milliseconds per step for the experiments of Fig. 2d

PRK1	KSL	KLS	AFE	PRH	APRH	PRK2	PRK3
5.27	5.41	5.88	7.71	12.62	17.38	10.41	14.16

standard numerical integration methods fail to provide a good approximation unless the step size is taken to be very small. Projector-splitting integrators for DLRA avoid such step size restrictions as the error convergence results are independent of the smallest non-zero singular value of the approximation. These schemes are commonly qualified as robust to small singular values. The robustness property was shown for the KSL scheme [18, Theorem 2.1] and the KLS scheme [8, Theorem 4]. The PRK methods also enjoy the robustness property [19, Theorem 6]. In the following, we experimentally study the robustness of the AFE, PRH and APRH integration schemes to the presence of small singular values.

The typical setting to assess the stability to small singular values of a given integration scheme, considers a matrix curve $t \in [0, T] \mapsto A(t) \in \mathbb{R}^{n \times n}$ of the form

$$A(t) = U(t)\Sigma(t)V(t)^\top \tag{5.2}$$

with

$$U(t) := \exp(t\Omega_U), \quad \Sigma(t) := \exp(t)D, \quad V(t) := \exp(t\Omega_V), \tag{5.3}$$

for some $n \times n$ skew-symmetric matrices Ω_U, Ω_V and a diagonal matrix $D = \text{diag}(\sigma_1, \dots, \sigma_n)$, for a positive and geometrically decaying sequence σ_i . A rank- r approximation of this curve is reconstructed by numerically integrating with the given scheme the DLRA Eq. (2.3) where the scalar field F is replaced by the exact derivative of the ambient curve (5.2):

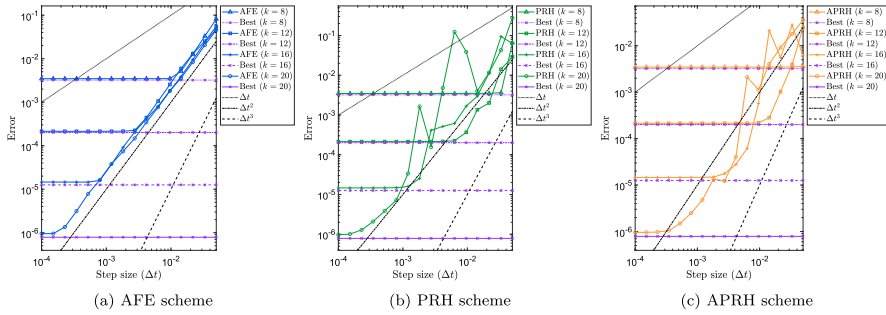


Fig. 4 Error at final time $\|Y_{\Delta t}(T) - A(T)\|_2$ versus the step size Δt of DLRA integration applied to (5.4) for reconstructing the curve (5.2) for different values of rank

$$A'(t) = U(t) \left(\Omega_U \Sigma(t) + \Sigma(t) + \Sigma(t) \Omega_V^\top \right) V(t)^\top. \tag{5.4}$$

The approximation error at final time is constituted mainly of the integration error which can be reduced by decreasing the step size, and the modeling error is affected only by the choice of r . A scheme is said to be robust to small singular values, if the integration error is independent of the choice of r . In practice, one must observe that the trend of the error as a function of the step size is unaffected by the choice of r for step sizes where the modeling error is negligible compared to the integration error.

Figure 4 shows the results for the experiment described in the previous paragraph with a curve of the form (5.2) with randomly generated Ω_U and Ω_V , initial singular values $\sigma_i = 2^{-i}$ and $n = 100$. The panels from left to right correspond respectively to the AFE, the PRH and the APRH schemes. Note that for the AFE and the APRH schemes, we use the exact expression for the second derivative of (5.2) given by

$$A''(t) = U(t) \left(\Omega_U^2 \Sigma(t) + \Sigma(t) + \Sigma(t) (\Omega_V^2)^\top + 2\Omega_U \Sigma(t) + 2\Omega_U \Sigma(t) \Omega_V^\top + 2\Sigma(t) \Omega_V^\top \right) V(t)^\top. \tag{5.5}$$

The results for AFE show the ideal outcome: the error curves for increasing values of r are superimposed until the modeling error plateau determined by the value of r is reached. These results empirically suggest that the AFE integration scheme is robust to small singular values. On the other hand, the PRH and APRH schemes which rely on retraction-based Hermite interpolation suffer from small singular values. Panels (b) and (c) of Fig. 4 exhibit the same oscillatory convergence trend that could be observed for both schemes in the experiments on the differential Lyapunov equation in Sect. 5.1.

A partial explanation for the oscillatory behavior observed in panels (b) and (c) of Fig. 4 for the PRH and the APRH schemes comes from studying robustness of the retraction-based Hermite interpolant (4.15) to the presence of small singular values at the interpolation points. Consider the following experiment. Take $Y_0(\sigma_r) \in \mathcal{M}_r \subset \mathbb{R}^{n \times n}$ with $n = 100, r = 12$ defined by

$$Y_0(\sigma_r) = U_0 \text{diag}(1, \dots, \sigma_r) V_0^\top$$

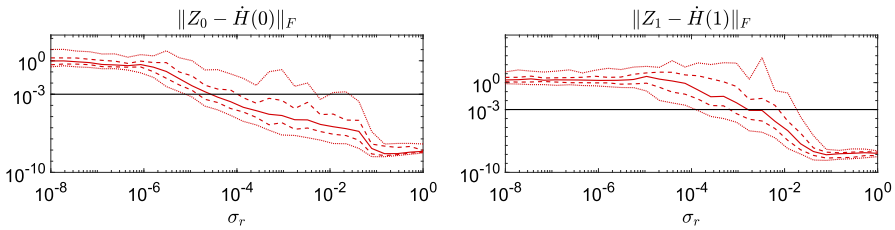


Fig. 5 Robustness to small singular values of the retraction-based Hermite interpolant (5.6). The solid line is the median of the error over 100 randomly generated instances for each value of σ_r while the dashed and dotted lines correspond to the percentiles [0.05, 0.25, 0.75, 0.95] of the sampled error

for some randomly generated orthogonal matrices U_0 and V_0 and with σ_i logarithmically spaced on the interval $[\sigma_r, 1]$, for some $\sigma_r \leq 1$. To obtain the second interpolation point, we first move away from Y_0 with the orthographic retraction along a random tangent vector $Z \in T_{Y_0}\mathcal{M}_r$ such that $\|Z\|_F = 1$ to get $\tilde{Y}_1 = R_{Y_0(\sigma_r)}(Z)$. Then, the second interpolation point is Y_1 , obtained from \tilde{Y}_1 by replacing its singular values with

$$\sigma_i(Y_1) = \sigma_i(Y_0)(1 + \xi_i),$$

for some random ξ_i drawn from a uniform distribution on $[1/2, 2]$. This way, the singular values decay of both Y_0 and Y_1 mimic a situation encountered in one step of the PRH and APRH integration schemes, when the smallest singular value of the current approximation is of the order of σ_r . Then, we randomly generate $Z_0 \in T_{Y_0}\mathcal{M}_r$ and $Z_1 \in T_{Y_1}\mathcal{M}_r$ with $\|Z_0\| = \|Z_1\| = 1$ and form the retraction-based interpolant (4.15) given by

$$H(\tau) = H(\tau; (0, Y_0, Z_0), (1, Y_1, Z_1)), \quad \tau \in [0, 1]. \tag{5.6}$$

For different values of the smallest singular value σ_r , we measure the discrepancies $\|Z_0 - \dot{H}(0)\|_F$ and $\|Z_1 - \dot{H}(1)\|_F$, where derivatives of H are obtained by finite differences. The experiment is repeated for each σ_r on 100 randomly generated instances and the error distribution is plotted against σ_r in Fig. 5. These results unequivocally indicate the fragility of retraction-based Hermite interpolant on the fixed-rank manifold when small singular values are present in the interpolation points. As σ_r decreases, the velocity error in $\tau = 0$ increases, and even more severely in $\tau = 1$. The fact that the error is non-negligible even for moderately small values of σ_r suggests that the PRH and APRH integration schemes may occasionally employ inaccurate retraction-based Hermite interpolants. This may contribute to the oscillatory behavior of the error observed PRH and APRH in Figs. 2 and 4.

6 Conclusion

This work contributes to strengthening the connection between retractions and numerical integration methods for manifold ODEs and especially DLRA techniques. In particular, we show that the so-called unconventional integration scheme [8] defines a

second-order retraction which approximates up to high-order terms the orthographic retraction. It remains an open question whether the same observation can be made for the recently proposed parallelized version of KLS [7].

We also derive three numerical integration schemes expressed in terms of retractions and showcase their performance on classic problem instances of DLRA. The derivation and the numerical results show that the methods can achieve second-order error convergence with respect to the time integration step. However, the methods have shown mixed results. While the AFE and the APRH schemes exhibit instability in the presence of large normal components of the ambient vector field, the PRH scheme appears more resilient to this aspect. On the other hand, the occurrence of small singular values in the approximation had no apparent effect on the performance of AFE but for the PRH and APRH methods, small singular values may explain occasional deviations from the expected second-order convergence behavior. We observe that the PRH scheme delivers similar performance, both with respect to computational time and accuracy, compared to its existing counterpart, PRK2. However, the high-order accelerated version, APRH, is found to be less favorable compared to analogous schemes such as PRK3. This is largely due to the additional computational cost incurred by the Weingarten map.

For other low-rank tensor formats, such as the Tucker or the tensor-train formats, retractions have also been proposed [21, 28]. However, to the best of our knowledge no retraction with an efficiently computable inverse retraction is known and the orthographic retraction has remained elusive due to the complexity of the normal space structure for these manifolds. Yet, the KLS scheme has been extended to low-rank tensor manifolds [8, §5]. Hence, assuming the connection with the orthographic retraction carries over to the tensor setting, it may be possible to retrieve the orthographic retraction for such tensor manifolds as a small perturbation of the KLS update. Then, the possibility to easily compute the inverse orthographic retraction would enable using also in the case of low-rank tensor manifolds the retraction-based endpoint curves and the numerical integration schemes presented in this work.

Funding Open access funding provided by EPFL Lausanne.

Declarations

Conflict of interest Not applicable.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Absil, P.A., Mahony, R., Sepulchre, R.: *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, Princeton (2008). <https://doi.org/10.1515/9781400830244>
2. Absil, P.A., Mahony, R., Trunpf, J.: An extrinsic look at the Riemannian Hessian. In: *Geometric science of information, Lecture Notes in Comput. Sci.*, vol. 8085, pp. 361–368. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-40020-9_39
3. Absil, P.A., Malick, J.: Projection-like retractions on matrix manifolds. *SIAM J. Optim.* **22**(1), 135–158 (2012). <https://doi.org/10.1137/100802529>
4. Absil, P.A., Oseledets, I.V.: Low-rank retractions: a survey and new results. *Comput. Optim. Appl.* **62**(1), 5–29 (2015). <https://doi.org/10.1007/s10589-014-9714-4>
5. Boumal, N.: *An Introduction to Optimization on Smooth Manifolds*. Cambridge University Press, Cambridge (2023). <https://doi.org/10.1017/9781009166164>
6. Boumal, N., Mishra, B., Absil, P.A., Sepulchre, R.: Manopt, a Matlab toolbox for optimization on manifolds. *J. Mach. Learn. Res.* **15**(42), 1455–1459 (2014)
7. Ceruti, G., Kusch, J., Lubich, C.: A parallel rank-adaptive integrator for dynamical low-rank approximation (2023). ArXiv preprint: [arXiv:2304.05660](https://arxiv.org/abs/2304.05660)
8. Ceruti, G., Lubich, C.: An unconventional robust integrator for dynamical low-rank approximation. *BIT* **62**(1), 23–44 (2022). <https://doi.org/10.1007/s10543-021-00873-0>
9. Charous, A., Lermusiaux, P.F.J.: Dynamically orthogonal Runge–Kutta schemes with perturbative retractions for the dynamical low-rank approximation. *SIAM J. Sci. Comput.* **45**(2), A872–A897 (2023). <https://doi.org/10.1137/21M1431229>
10. Einkemmer, L., Lubich, C.: A low-rank projector-splitting integrator for the Vlasov–Poisson equation. *SIAM J. Sci. Comput.* **40**(5), B1330–B1360 (2018). <https://doi.org/10.1137/18M116383X>
11. Feppon, F., Lermusiaux, P.F.J.: A geometric approach to dynamical model order reduction. *SIAM J. Matrix Anal. Appl.* **39**(1), 510–538 (2018). <https://doi.org/10.1137/16M1095202>
12. Hairer, E., Lubich, C., Wanner, G.: *Geometric numerical integration, Springer Series in Computational Mathematics*, vol. 31. Springer, Heidelberg (2010). *Structure-preserving algorithms for ordinary differential equations, Reprint of the second edition* (2006)
13. Hairer, E., Nørsett, S.P., Wanner, G.: *Solving ordinary differential equations. I, Springer Series in Computational Mathematics*, vol. 8, second edition. Springer-Verlag, Berlin (1993)
14. Hairer, E., Wanner, G.: *Solving ordinary differential equations. II, Springer Series in Computational Mathematics*, vol. 14, revised edn. Springer-Verlag, Berlin (2010). <https://doi.org/10.1007/978-3-642-05221-7>
15. Helmke, U., Moore, J.B.: *Optimization and dynamical systems. Communications and Control Engineering Series*. Springer-Verlag London, Ltd., London (1994). <https://doi.org/10.1007/978-1-4471-3467-1>
16. Jahnke, T., Huisinga, W.: A dynamical low-rank approach to the chemical master equation. *Bull. Math. Biol.* **70**(8), 2283–2302 (2008). <https://doi.org/10.1007/s11538-008-9346-x>
17. Kazashi, Y., Nobile, F., Vidličková, E.: Stability properties of a projector-splitting scheme for dynamical low rank approximation of random parabolic equations. *Numer. Math.* **149**(4), 973–1024 (2021). <https://doi.org/10.1007/s00211-021-01241-4>
18. Kieri, E., Lubich, C., Walach, H.: Discretized dynamical low-rank approximation in the presence of small singular values. *SIAM J. Numer. Anal.* **54**(2), 1020–1038 (2016). <https://doi.org/10.1137/15M1026791>
19. Kieri, E., Vandereycken, B.: Projection methods for dynamical low-rank approximation of high-dimensional problems. *Comput. Methods Appl. Math.* **19**(1), 73–92 (2019). <https://doi.org/10.1515/cmam-2018-0029>
20. Koch, O., Lubich, C.: Dynamical low-rank approximation. *SIAM J. Matrix Anal. Appl.* **29**(2), 434–454 (2007). <https://doi.org/10.1137/050639703>
21. Kressner, D., Steinlechner, M., Vandereycken, B.: Low-rank tensor completion by Riemannian optimization. *BIT* **54**(2), 447–468 (2014). <https://doi.org/10.1007/s10543-013-0455-z>
22. Kusch, J., Ceruti, G., Einkemmer, L., Frank, M.: Dynamical low-rank approximation for Burgers’ equation with uncertainty. *Int. J. Uncertain. Quantif.* **12**(5), 1–21 (2022). <https://doi.org/10.1615/int.j.uncertaintyquantification.2022039345>
23. Lee, J.M.: *Introduction to Riemannian manifolds, Graduate Texts in Mathematics*, vol. 176. Springer, Cham (2018)

24. Lubich, C., Oseledets, I.V.: A projector-splitting integrator for dynamical low-rank approximation. *BIT* **54**(1), 171–188 (2014). <https://doi.org/10.1007/s10543-013-0454-0>
25. Peng, Z., McClarren, R.G., Frank, M.: A low-rank method for two-dimensional time-dependent radiation transport calculations. *J. Comput. Phys.* **421**, 109735 (2020). <https://doi.org/10.1016/j.jcp.2020.109735>
26. Séguin, A., Kressner, D.: Hermite interpolation with retractions on manifolds (2022). ArXiv preprint: [arXiv:2212.12259](https://arxiv.org/abs/2212.12259)
27. Shub, M.: Some remarks on dynamical systems and numerical analysis. In: *Dynamical systems and partial differential equations* (Caracas, 1984), pp. 69–91. Univ. Simon Bolivar, Caracas (1986)
28. Steinlechner, M.: Riemannian optimization for high-dimensional tensor completion. *SIAM J. Sci. Comput.* **38**(5), S461–S484 (2016). <https://doi.org/10.1137/15M1010506>
29. Séguin, A.: Retraction-based numerical methods for continuation, interpolation and time integration on manifolds (2023). PhD Thesis, EPFL
30. Uschmajew, A., Vandereycken, B.: Geometric methods on low-rank matrix and tensor manifolds. In: *Handbook of Variational Methods for Nonlinear Geometric Data*, pp. 261–313. Springer, Cham ([2020] 2020). https://doi.org/10.1007/978-3-030-31351-7_9
31. Vandereycken, B.: Low-rank matrix completion by Riemannian optimization. *SIAM J. Optim.* **23**(2), 1214–1236 (2013). <https://doi.org/10.1137/110845768>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.